

(12) UK Patent Application (19) GB (11) 2 282 018 (13) A

(43) Date of A Publication 22.03.1995

(21) Application No 9318979.3

(22) Date of Filing 14.09.1993

(71) Applicant(s)
Northern Telecom Limited

(Incorporated in Canada - Quebec)

World Trade Center Of Montreal,
380 St Antoine Street West, 8th Floor, Montreal,
Quebec H2Y 3Y4, Canada

(72) Inventor(s)
John Edwin Midwinter
Martin Sabry

(74) Agent and/or Address for Service
J P W Ryan
Northern Telecom Europe Limited, Patents and
Licensing, West Road, HARLOW, Essex, CM20 2SH,
United Kingdom

(51) INT CL⁶

H04B 10/12 10/13 10/20, H04J 4/00 14/00, H04Q
11/00

(52) UK CL (Edition N)
H4B BK12M BK20

(56) Documents Cited

GB 2184327 A

EP 0497670 A

EP 0463634 A

EP 0400959 A

EP 0351729 A

EP 0268355 A


(58) Field of Search

ONLINE DATABASES: WPI

(54) Optical communications network

(57) An optical communications system comprises an inner core mesh network, 13, a plurality of outer unidirectional ring networks 14 coupled to the core, and passive optical local networks 12 coupled to the outer networks whereby to provide access to terminals 11 coupled to the network. Information is transported across the network using time division multiplexing in combination with frequency division multiplexing; capacity is allocated to particular routes by giving those routes certain frequency/time-slot-within-multiplexing-cycle combinations. A line of the network may comprise a plurality of optical fibres.

Fig.1.



OPTICAL COMMUNICATIONS NETWORK

This invention relates to communications networks and in particular to optical communications networks.

Optical fibre transmission is a widely used technique in the communications field. The technique has the major advantages of low attenuation and high bandwidth in comparison with electronic transmission techniques. However, in communications network applications, the use of fibre has not extended significantly beyond trunk transmission. Thus, local transmission and switching are performed electrically. This requirement for both electrical and optical communications introduces two significant limitations. Firstly there is the need to provide electronic/optical and optical/electronic interfaces between the two transmission media. Secondly, the bandwidth restrictions of the electrical part of the network present the full potential bandwidth available in the optical spectrum from being utilised. For example, the complete radio and microwave spectrum has a bandwidth of about 300GHz, whereas the potential bandwidth available in the optical spectrum in the typical transmission wavelength range of 1500 to 1600 nm is of the order of 4000 to 6700 GHz.

Furthermore, in a communications network, it is necessary to provide a number of levels of multiplexing, the highest level being used for trunk transmission. This necessitates the provision of multiplexing and demultiplexing equipment. Although the recent introduction of synchronous systems (SDH) to replace the present plesiochronous systems (PDH) has resulted in a reduction of the volume of multiplexing/demultiplexing equipment required, this

equipment still represents a significant proportion of the overall system cost.

It is an object of the invention to minimise or to overcome those disadvantages.

According to the invention there is provided a communications network in which information is transported optically between terminals coupled to the network, and wherein said information is carried in pixels in a discrete communications space defined by time and wavelength co-ordinates.

According to the invention there is further provided an optical communications network, including an inner core network having a plurality of nodes, outer networks coupled to at least some of said nodes, and local distribution networks each coupled to a said outer network and each providing access to a plurality of terminals, wherein information transported via the network between terminals is carried in pixels in a discrete communications space defined by time and wavelength coordinates.

The communications space extends throughout the network

The network architecture is designed to be integrated with, and eventually replace, the current network in incremental steps from the top down using a common set of components throughout. In general this will ensure that maximum benefit can be derived by spreading costs amongst users at the higher network levels, and reductions in component costs, due to mass production of standardised modules at the subsequent lower levels.

In the first stage of evolution we envisage existing transmission links within an inner core network being upgraded to use Dense WDM transmission but with each carrier operating within the Synchronous Digital Hierarchy (SDH) to standard interfaces. Later these interfaces would be replaced by transparent wavelength switches and thereby facilitate transmission between rings entirely in the optical domain. Ultimately, individual customers could be serviced

directly through passive optical networks (PONs) optically connected to the rings.

In a high capacity network carrying traffic multiplexed from many users, one would not expect the mean traffic flow along routes to fluctuate by large amounts in short times. Accordingly, we propose that the traffic carrying ability in the inner and outer core networks (mesh and rings) should be allocated on a slowly time varying basis (~hours), initially in blocks of one optical carrier e.g. at the synchronous transport module standard STM-4. The choice of STM-4 may seem low, at a time when much interest centres on transmitting at STM-16 or higher rates, however, we have two reasons for proposing it. Firstly, a nation-wide transparent optical network with transmission at 1550nm over distances up to the order of 1000km is most likely to be limited to a data rate less than 1Gbit/s per carrier due to material dispersion, as much of the dark fibre installed and the fibre already used is optimised for transmission around 1300nm. Secondly, in order to minimise the amount of time domain processing on the multiplexed signals we are concerned to identify a traffic capacity building block appropriate as the minimum unit of capacity to be allocated between any two core nodes in the network. However, as the network is extended to include rings and the opto-electronic boundary is brought closer to the customers, we can introduce a limited degree of time division (albeit without time shifting buffers at nodes) in order to further sub-divide the communication space. Recognising that telecommunications is based around an 8kHz sampling cycle corresponding to a 125µs frame duration, we propose dividing it into 16 sub-frames (SF), each of about 7.8µs. Thus the minimum capacity block would be reduced to about 38Mbit/s. By using more sophisticated technology the basic building block units may for example be designed around the higher rate STM-16. It will be appreciated that the technique is in no way limited to any particular rate. As the time/wavelength communications space extends throughout the network. This obviates the need for any multiplexing or demultiplexing facilities thus providing a significant cost saving in comparison with a conventional network

An embodiment of the invention will now be described with reference to the accompanying drawings, in which:-

Fig. 1 is a schematic diagram of a communications network according to the invention;

Fig. 2 illustrates the principle of communication within the network of Fig. 1;

Fig. 3 illustrates the communication process at a core node of the network of Fig. 1;

Fig. 4 shows the schematic construction of a core node for the network of Fig. 1;

Fig. 5 shows the schematic construction of a ring node of the network of Fig. 1;

Fig. 6 shows an interface between the network and a user terminal;

Fig. 7 shows a core network traffic module;

Fig. 8 illustrates the operation of a ring-passive optical network interface; and

Figs. 9, 10 and 11 illustrate mean success rate and throughput from the network of Fig. 1;

Referring to Fig. 1, the network provides communication between a plurality of terminals 11 via access passive optical networks (PON's) 12 and a network core comprising an inner mesh core 13 and outer core rings 14. The inner core 13, which comprises the highest hierarchical level of the network, is a highly interconnected mesh. Each of the core nodes 15 provides routing within the inner core and at least some core nodes also provide access to respective outer core rings 14. The latter are provided with ring nodes 16 which nodes access the passive optical networks 12 to which the user terminals 11 are coupled. In the arrangement of Fig. 1, traffic is carried end-to-end between terminals in a wholly optical manner using time division and wavelength division multiplexing (WDM).

Fig. 2 illustrates the way in which information is transported in the network of Fig. 1. The figure illustrates the concept of a communications space within an optical fibre transmission path. Information is carried in pixels 21 in a two dimensional discrete

communications space 20, each pixel being defined by respective wavelength (λ) and time (t) co-ordinates. The communications space is the same in any fibre of the network whether that fibre is part of a trunk line or is a user link to a subscriber terminal. The shaded plane 20 in Fig. 2 represents the communication space in a single fibre. As has already been pointed out, it is possible to wavelength multiplex and one can therefore think of the frequency dimension as divided into bands each representing a transparent optical channel. In addition, time is also partitioned analogous to conventional Time Division Multiplexing (TDM). Combined, the two multiplexing techniques result in a matrix of pixels each of which may be described by a unique (λ, T) reference within the fibre. Thus each pixel within the communication space might correspond to one optical STM (synchronous transport module) channel. A Third dimension may be added to accommodate the possibility that more than one fibre may exist between any two points in a network. In contrast to the dimensions of the time-channel plane, which are the same at any point in any optical network, the extent of the fibre dimension may vary through a network. Moreover, in order to maintain all optical transmission paths the only type of operation that can be permitted when switching is to shift the time-wavelength plane (or parts of it, say one pixel) within this dimension, i.e., space switching.

A node in a mesh network will generally have several input and output fibres connecting it to a number of different nodes. Given that nodes have the ability to separate the individual wavelength slots on each of their input fibres and reassemble them in a different configuration on their output fibres, we can use the fact that the output fibres lead to different nodes, to route wavelength channels through the network. If we further assume that the amount of time required by the node to re-configure the mapping of wavelength slots is relatively small compared to the degree of time division, each of the pixels in Fig. 2 may be thought of as a "temporary" transparent data channel which can be used to make an arbitrary point to point connection in the network. As we are proposing all optical switching, the pixels must maintain their position within the time-wavelength planes of the fibres, from the input to the output of

a node and indeed throughout the whole multinode network. Fig. 3 shows a hypothetical all optical mapping for a node with 3 input and 3 output fibres each filled (indicated by the shaded pixels) to half of the theoretically available capacity. To perform switching entirely within the optical domain, each pixel maintains its position within the time-wavelength plane of the communications space. Accordingly the pixels are mapped by shifting them along the fibre axis

The functionality of a core node is illustrated in a different manner in Fig. 4. The node comprises a wavelength switch, without wavelength or time shifting, which controls paths between core and ring inputs 41, 42 and core and ring outputs 43, 44 respectively.

Switching between rings and core nodes, as well as from a ring to a core and vice-versa are all required. However, as the capacity in both the core and the rings is allocated in subframes, if M is the total number of fibres entering (and leaving) each node, w the number of active wavelength channels per fibre and F the number of sub-frames at the basic building block rate of 38MBit/s (or 155MBit/s), then the node can be seen to be a switch receiving $M \cdot w \cdot F$ input channels and cross connecting them to $M \cdot w \cdot F$ output channels, irrespective of whether the individual fibres are part of a ring or the mesh network. In other words, a single type of core node may be provided, and it may be used either in conjunction with rings or purely as part of the inner core network switching fabric. As a consequence the rings should be integer multiples of sub-frames in circumference as well as uni-directional, in order to maintain both the timing within the inner core network and their self-healing properties. In a typical arrangement, the total number M of fibres may be 10, the number W of wavelength channels may be 50 and the number F of subframes may be 16 giving a throughput of between 0.3 and 1.24TBit/s made up of 8000 independently routed basic channels.

In order to synchronise the switching of subframes in the core network, all sections of fibre in that network should be of a length corresponding to integral multiples of the subframe interval. For example, an 8kHz frame rate divides into 16 subframes implies, on

average, an additional 0.75 km of fibre in each cable termination. Moreover, a general band of about 700 ns between data contained in consecutive subframes is desirable to allow for differences in propagation delay caused by dispersion effects.

Although capacity may be allocated in blocks of 38MBit/s (or 155MBit/s) within and between rings, if the opto-electronic border is to be brought out to the individual customers, the network must realistically be capable of providing as little capacity as 64kBit/s channels. We therefore propose that the sub-frames be composed of time slots of a duration sufficient for the transmission of standard ATM (asynchronous transfer mode) cells (about 680 ns at STM-4 or about 170 ns at STM-16) which may be dropped into or extracted from a PON individually. As a result, it is possible to transmit information from customer to customer entirely within the optical domain in denominations ranging from an ATM cell upwards.

The ring-PON interface is shown in Fig. 5. Similar to the core node, it accesses all the incoming wavelength channels. Several input and output fibres may exist on the ring level. However, as all these fibre by definition pass through the same sequence of nodes, no subframe to subframe switching is required. The real purpose of these switches is to establish access to one or more PONs from the ring. Consequently, the routing which takes place is based on a sub-frame to ATM cell (C) transformation. The nodes continuously examine optical pre-headers which preceded every ATM cell (indicating the cells' final destination PON on the ring) and then direct them individually, either to the appropriate PON, or to the ring level output fibre for further transmission down the line of PONs. Depending on the exact configuration of the ring-PON interface the pre-header may be lost when the ATM cell is dropped into a PON, but as no further routing decisions are required before the customer terminal this is of no consequence. In the cases where a ring is made up of several fibres, the node also ensures that packets injected from the customers on a PON are directed to the correct ring fibre. Connections between customers on different rings are then established by injecting ATM cells into vacant slots in a sub-frame channel that connects the relevant rings. Since the routing

pattern of the sub-frames is predetermined at the inner core level, the data channels i.e. the pixels within the fibre communication space may be viewed as set of postal pigeon holes each addressing a different destination ring. Consequently, the transmission of information from source to destination is reduced to the insertion and removal of ATM cells at the two involved ring-PON interfaces. Moreover, as these are both strictly localised processes, maintaining precise timing is required within the individual PONs.

Communication between customers on the same ring is achieved by allocating additional sub-frame channels used for local traffic only. If the two customers happen to be on the same PON or on PONs connected to the ring at the same point, the ATM cells simply make one complete round trip before they are dropped back into the PON. Alternatively, an additional level of switching may be introduced, indicated by the dashed lines in Fig. 5 that would allow ATM cells belonging to this category of traffic to be routed directly within the ring-PON node. This would not require the introduction of a new level of functionality, as the processing necessary is a subset of that, which has already been specified for the other kinds of traffic managed by this type of node.

It will be appreciated that electronic interfaces will be required at the periphery of the network. In line with the development strategy outlined above, the specification of these opto-electronic interfaces must remain the same throughout the expansion of the network. As shown in Fig. 6 the optics consist of two main components; a tuneable narrow line-width laser TX and a photo-detector RX. A tuneable wavelength discriminatory device may also be inserted into the receiving path if one desires the terminal to be able to single out a unique wavelength channel if several appear on the input fibre simultaneously. The two electronic modules are the control circuitry CTRL required for the correct operation of the terminal and high speed buffer memories. BUFR and BUFT are associated with the photodetector and the laser respectively. The buffer memory BUFR ensures that when an ATM cell destined for the terminal is detected by the optical receiver it may be stored temporarily in order to allow the data to be read out at the lower bit rate anticipated on the

electronic side of the interface. In this way the only high speed electronics (STM-4 or STM-16) used in the network are the two 53 byte memory cells in each terminal.

There now follows an analysis of the communications network described above. This analysis is given by way of example and demonstrates the flexibility of the techniques described herein.

Analysis of the core network

As previously stated, the function of the core network is to establish high capacity "transparent" data channels between geographically localised communities of interest. We assume the traffic in the inner core can be partitioned into R geographically localised communities, each of which can be covered by a ring, and that all have similar statistical characteristics in terms of local and inter community traffic flows. If the distribution of traffic originating from a ring is plotted versus the remaining rings, where these are ordered in sequence according to the proportion of the traffic they receive, we further assume this to show a substantially linear distribution with an order of magnitude difference between the most and least popular communities of interest as illustrated in the traffic module shown in Fig. 7. Consequently, if wqx is the number of channels available for external traffic in a ring, where w is the number of active wavelength channels per fibre, q is the number of files and F is the number of subframes, then the value of this packet must be maintained greater than or equal to $(10/2) \times R$. For example, if there is a minimum of 2 fibres in the rings, each with an equivalent of 50 wavelengths, each partitioned onto 16 sub-frames dedicated to core traffic, the network can accommodate a minimum of $50 \times 2 \times 16 = 320$ rings. Within the context of a national UK network this would imply rings with an average radius of roughly 30km assuming a uniform population density across the country. Alternatively, with 25×10^6 network subscribers, it implies ~78,000 customers per ring.

Although we envisage that only part of the time-wavelength-fibre space will be used for transmission at any one time, we also assumed that the connection pattern in the core would be (albeit slowly) varying dynamically. Consequently, there will be a finite

probability that a new channel along a given route can not be accommodated and therefore must be denied if requested. If the network control algorithms do not possess any "intelligence" and requests for new channels are made, as well as allocated, utterly at random, the highest probability for blocking, P_B , arises when trying to make a connection along the longest route where the transmission space, on all stages, is one channel short of being utilised to its maximum permitted level. If f_i denotes the maximum fraction of the total number of wavelength slots allowed to be active in each fibre, and q_i the number of fibres in the i^{th} cable along the route, then P_B is given by eqn. (1) where p is number of wavelength slots tested.

$$P_B = \left[1 - (1 - f_1^{q_1}) \times \dots \times (1 - f_{i-1}^{q_{i-1}}) \times (1 - f^{q_i}) \right]^p \quad (1)$$

Assuming that the filling factor, f and the number of the fibres, q , is the same throughout the whole of a route for a given connection and all vacant channels are tested, eqn. (1) may be rewritten as eqn (2) below, where q , l and f are defined as above and W is the number of channel slots in a fibre.

$$P_B = \left[1 - (1 - f^q)^l \right]^{(l-f \cdot q \cdot W)} \quad (2)$$

For example, if there are 8 and 2 fibres per cable in the core and ring networks respectively, with 100 wavelength slots and a 50% filling factor in each fibre, then the probability for not making a connection between two rings separated by 11 core nodes (longest path in 7x7 rectangular mesh) should be 5×10^{-28} on any particular chosen route.

Analysis of the access network

In addition to the blocking which may take place within core network, there exist an additional set of blocking probabilities within the access network which arise from contention between ATM cells. Fig. 8 depicts a conceptual view of the interaction between a ring consisting of two fibres and a PON. The upper part of the figure shows a hypothetical switching pattern for the ring-PON interface, depicted in a manner similar to that previously presented for the

core node, Fig. 3. The pixels resulting from the original partition of the communication space, i.e. the basic core communication channels are delimited by the bold lines. In accordance with the proposed architecture, these have been further divided into smaller areas within the time dimension to represent the individual ATM cells. As was the case for the capacity available in the core network, these carry a varying amount of traffic (indicated by the shaded areas) depending on the offered load and the existing traffic patterns. As previously stated, a ring-PON interface, represented in Fig. 5, does not perform any SF→SF switching, nor does it facilitate ATM cell to ATM cell switching between the ring input and output fibres.

In general there can exist three different types of contention within the access network. When an ATM cell is injected into the network by a terminal it will cause a collision if the given ATM slot already contains another cell. Accordingly, the network control should keep track of which of the slots in the various data channels in a ring are in use at any given point in time and only make those that are empty available for "new" connections. This implies that before a connection can be established between two terminals, the one which initiates the request must signal the network to determine the time and wavelength references of the appropriate channels which are available as well as which ATM slots are unused within them. Given that a ring consists of two or more fibres the control algorithm must also verify that two terminals in the same PON are not competing for identical ATM slots on different fibres as this will result in the cells colliding within the PON. In such cases, only one of the available slots may be allocated. The control issues associated with these two types of collision may be administered by the local hardware which oversees the dialogues between the PONs (terminals) and their rings. These localised control centres would though not store sufficient information about the whole network for a terminal to determine, at the time of initiating a call, whether the receiving terminal is available or in fact if the time-wavelength reference chosen for the transmission of ATM cells will cause contention with already existing connections between other PONs in other rings and the receiving PON. Thus, the third type of

contention arises from potential collisions in the receiving PON. Viewed from the position of a ring receiving a number of calls (or request for calls), however, this new set of administrative control issues are again localised. We therefore suggest that this network architecture lends itself to distributed control concentrated, for example, in each ring. Communication between terminals would then only require the involvement of the control centres directly affected, i.e. generally only two.

In order to further quantify the issues raised above, we present a general analysis of the various levels of contention and blocking identified in the access network. If x ATM cells are selected at random from $W.f.q$ ATM slots the probability that they are all destined for the same PON is given by P_x defined by eqn. (3).

$$P_x = \frac{1}{n} \left(1 - \frac{1}{n}\right)^{w.f.q-1} C_x^{w.f.q} \quad (3)$$

For example, if a ring has the dimensions used above and the ATM cells transmitted in a certain slot are distributed between 500 PONs (which in the context of a network with 25×10^6 customers and 320 rings implies ~150 customers per PON), the likelihood that a PON does not receive an ATM cell is $P_0 = 0.819$.

If terminals grab empty slots at complete random, the probability that a given slot is duplicated in a PON is given by eqn. (4). Naturally, this is impossible unless there are two or more fibres in a ring, i.e. for $q \geq 2$.

$$\frac{f.(q-1)}{N} \quad (4)$$

In the case where receivers are not wavelength selective, it follows that only one terminal per PON can receive an ATM cell per ATM slot. Consequently, in order to fully use the allocated transmission capacity, the number of PONs, N , in a ring must be greater than or equal to the number of active wavelengths. $N \geq W.f.q$. Given that one or more channels have been established between two rings the probability of making an ATM connection is the product of a slot not

being duplicated in the transmitting PON, and being empty in the receiving PONs, eqn. (5), where p is the number of empty ATM slots tested.

$$\left[1 - \left[\frac{1}{n} \cdot f \cdot (q - 1)\right]^r\right] \left[1 - [1 - p_n]^r\right] \quad (5)$$

For example, if two rings have 500 PONs and 2x50 active channels each, and only one 38MBit/s channel is connecting them, a single 64kBit/s line may on average be established 96.7% of the time, as long as the channel is operated at less than 99.6% of its total capacity. Alternatively, if the same rings only had 100 or 50 PONs, the number of customers per PON would increase to approximately 750 and 1500 and the probability for making a given connection would be reduced to 59.8% or 24.8% respectively. It is understood that, if we measure hardware in terms of the number of PONs (and hence size and number of ring-PON interfaces) and performance in terms of the likelihood of being able to make a given connection, we can generally trade one for the other. However, in order to determine the overall blocking which takes place in a ring it is necessary to take the traffic arriving from all the other rings in the whole network into account. Eqn. (6), and Fig. 9 express the probability of being able to allocate all the receiving capacity required ($W.f.q.f_R$) without contention (for $N=500$), given that the traffic pattern is totally random, where f_R is the fraction of the allocated capacity that is used. They indicate that it is extremely unlikely that it will be possible to fully allocate all the available capacity within a ring such that no contention exists, but that this should continuously become more readily achievable as the offered load is reduced.

$$\frac{N!}{[N - W.f.q.f_R]! N^{W.f.q.f_R}} \quad (6)$$

The average fraction \emptyset of injected ATM cells which are successfully transmitted, (evaluated as $1 - P_{\text{collision}}$) versus the ring traffic load can be expressed as the fraction of PONs receiving exactly one wavelength with respect to the mean number of filled ATM slots,

(mean number of PONs receiving exactly one ATM cell per ATM slot)/(W.f.q.fR), eqn (7).

$$\phi = \frac{\left[\frac{1}{n}\right]^n \left[1 - \frac{1}{n}\right]^{W \cdot f \cdot q \cdot f_R - 1} C_i^{W \cdot f \cdot q \cdot f_R \times N}}{W \cdot f \cdot q \cdot f_R} = \left[1 - \frac{1}{N}\right]^{W \cdot f \cdot q \cdot f_R - 1} \quad (7)$$

Fig. 10 plots ϕ and the gross throughput, given as the product of the load and the success rate, versus f_R . It reveals that rings with the dimensions used in the above examples, even without wavelength selective receivers, would have an 82% success (or 18% failure) rate, as well as throughput, when fully loaded. In general, the throughput curves calculated for the examples with 500 and 100 PONs per ring, do not exhibit maxima. This implies that the throughput will increase continuously in step with the load. In other words, we would not expect it to be necessary to build safeguards into the network control to cope with congestion as is currently proving to be necessary with conventional ATM. However, in the example where the ring only had 50 PONs the throughput shows a shallow maximum of about 0.5. This result suggests that when the number of PONs in a ring is small compared to the address space desired in the inner core (which is a function of the total number of rings and the respective traffic distributions) care should be taken to limit the injected traffic to the optimum level or the receivers in the terminals should be specified as wavelength selective.

If terminals are given the ability to select the individual wavelength channels appearing simultaneously on the input fibre, we can remove the restriction of a single ATM cell per PON per ATM slot. In its place, it will be necessary to prevent two ATM cells being transmitted in identical ATM slots on different fibres (presupposes $q \geq 2$), being dropped into the same PON at the same time. Since the W.f active channels are selected at random from W in each of q fibres the probability that a given wavelength channel is used exactly j times is given as Q_j .

$$Q_j = f^j [1 - f]^{q \cdot \text{ring} - j} C^{q \cdot \text{ring}} \quad (8)$$

The probability that this wavelength is chosen more than 1 time if x wavelengths are chosen out of $W.f.q$ is:

$$P_{j,x} = 1 - \frac{C_o^j \cdot C_x^{W.f.q-j} + C_l^j \cdot C_{r-l}^{W.f.q-j}}{C_x^{W.f.q-j}} \quad (9)$$

Hence, the mean number of packets that collide in a PON per ATM slot is PQ .

$$PQ = \sum_{x=2}^{W.f.q} \left(P_x \cdot W \cdot \sum_{j=2}^q Q_j \cdot P_{j,x} \right) \quad (10)$$

Fig. 11 plots the expected success rate and gross throughput for a ring where all terminals have wavelength selective receivers, eqn. (11), versus the offered load (where f in PQ is replaced by $f \cdot f_R$).

$$\text{Throughput} = 1 - \frac{PQ \cdot N}{W \cdot f \cdot q \cdot f_R} \quad (11)$$

A comparison between Fig. 10 & 11 reveals that this modification of the customer terminals causes a marked improvement in all three examples. The decision of whether to introduce the additional complexity and cost of wavelength selectivity in each terminal will ultimately depend upon the required performance and the relative dimensions of the inner and the outer core, and the number of PONs in each ring.

The throughput graphs, eqn. (7 & 11), plotted in Figs. 10 & 11 represent the total fraction of traffic passing through the network and arriving uncorrupted at its destination. Accordingly $(1 - \text{Throughput})$ describes the proportion of data which is lost or corrupted. Therefore, depending on the protocol, this may be equated to either a Bit Error Rate (BER), or to the proportion of the gross offered load which is retransmitted data. If the basic protocol does not incorporate retransmission, the net throughput will be equal to the gross throughput but with a Packet Error Rate imposed by the protocol. In the case where the terminals are wavelength selective this equates to a BER ranging from 5×10^{-3} to 5×10^{-4} for the three examples given. In comparison, if all corrupted packets are retransmitted the control itself will not give rise to a BER, although

the net throughput will be reduced by a factor equal to the original success rate, i.e. $\text{net throughput} = \text{offered load} \times \text{success rate}$. For example, if the overall transmission rate is 2.4 GBit/s (corresponding to STM.16) the aggregate capacity per instruction at 100% loading is 2 MBit/s for non-wavelength selective receivers and 3 MBit/s for wavelength selective receivers. So, similarly to the hardware vs. performance trade-off described above, we find that there exists a relationship between protocol issues and the performance/characteristics of the network.

CLAIMS:

1. A communications network in which information is transported optically between terminals coupled to the network, and wherein said information is carried in pixels in a discrete communications space defined by time and wavelength coordinates.
2. An optical communications network, including an inner core network having a plurality of nodes, outer networks coupled to at least some of said nodes, and local distribution networks each coupled to a said outer network and each providing access to a plurality of terminals, wherein information transported via the network between terminals is carried in pixels in a discrete communications space defined by time and wavelength coordinates.
3. An optical network as claimed in claim 2, wherein each said pixel corresponds to an optical synchronous transport module channel.
4. An optical network as claimed in claim 2 or claim 3, wherein switching is effected by space switching of pixels.
5. An optical network as claimed in any one of claims 2 to 4, wherein said inner core network comprises a mesh.
6. An optical network as claimed in claim 5, wherein each said outer network comprises a ring.
7. An optical network as claimed in claim 6, wherein said ring is unidirectional.
8. An optical network as claimed in claims 5, 6 or 7 wherein each said local network comprises a passive optical network (PON).
9. An optical network as claimed in claim 6 or 7, wherein capacity within the inner core and the outer network rings is

allocated in subframes, and wherein each said ring corresponds in circumference to an integral number of subframes.

10. An optical communications network substantially as described herein with reference to and as shown in the accompanying drawings.

11. A method of optical transmission of information in a communications network, the method comprising transporting the information in pixels, in a discrete communications space defined by time and wavelength coordinates.

12. A method of optical transmission of information substantially as described herein with reference to and as shown in the accompanying drawings.

19

Relevant Technical Fields

(i) UK Cl (Ed.)

(ii) Int Cl (Ed.)

Databases (see below)

(i) UK Patent Office collections of GB, EP, WO and US patent specifications.

(ii) ONLINE DATABASE : WPI

Search Examiner
DR E P PLUMMER

Date of completion of Search
10 NOVEMBER 1993

Documents considered relevant
following a search in respect of
Claims :-
ALL

Categories of documents

- | | |
|--|---|
| <p>X: Document indicating lack of novelty or of inventive step.</p> <p>Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.</p> <p>A: Document indicating technological background and/or state of the art.</p> | <p>P: Document published on or after the declared priority date but before the filing date of the present application.</p> <p>E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.</p> <p>&: Member of the same patent family; corresponding document.</p> |
|--|---|

Category	Identity of document and relevant passages		Relevant to claim(s)
X	GB 2184327 A	(ATT) - eg abstract, Figure 2	1, 2, 11 at least
X	EP 0497670 A	(ALCATEL) - whole document	1, 11
X	EP 0463634 A	(NEC) - eg Claim 1	1, 11
X	EP 0400959 A	(TOSHIBA) - eg Figure 1, Claim 1, column 8, line 44 ff	1, 2, 11 at least
X	EP 0351729 A	(FUJITSU) - eg figure 16	1, 11
X	EP 0268355 A	(HICKS) - eg abstract	1, 11

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).

Fig.1.




Fig.2.






Fig.3.



2/8



3/8

Fig.4.




Fig.5.




Fig.6.






Fig.9.



$W^* f^* q^* F = (10/2)^* R$
 $W = 100$
 $f = 0.5$
 $q = 2$
 $F = 16$
 $R = 320$



GEOGRAPHICALLY LOCALISED COMMUNITIES OF INTEREST

Fig. 7.





Fig. 8a







Fig.8b





8/8

Fig.10.




Fig.11.

